Patents and the Performance of Voluntary Standard Setting Organizations *

Marc Rysman
Boston University, Department of Economics

Tim Simcoe
J.L. Rotman School of Management, University of Toronto

June 2006

*Financial support for this research was provided by CITRIS, and the NET Institute. Useful comments were received from Kevin Lang, Josh Lerner, David Mowery, Bronwyn Hall, Avi Goldfarb, Shane Greenstein, Ken Corts, Katrin Cremers, Michael Ward and seminar participants at Brandies University, UC Berkeley, the International Industrial Organization Conference in Boston (2006) and the NET Institute Conference (NY, 2006). Address for correspondence: Joseph L. Rotman School of Management, 105 St. George Street, Toronto, ON M5S 3E6,Canada. E-mail: timothy.simcoe@rotman.utoronto.ca

Patents and the Performance of Voluntary Standard Setting Organizations

Abstract

This paper examines the economic and technological significance of voluntary standard setting organizations (SSOs). These groups are common in industries with strong network effects, providing a forum for collective decision-making and an alternative to coordination through market competition or government regulation. We use patent citations as a measure of SSO performance. Specifically, we model the flow of citations to a sample of U.S. patents disclosed during the standard-setting process at four major SSOs. Our main results show that the age distribution of SSO patent citations is skewed towards later years (relative to an average patent), and that citations increase substantially following disclosure. This suggests that SSOs not only identify and select appropriate technologies, but also play an important role in promoting the adoption of those technologies as industry standards. These results provide the first empirical look at patents disclosed to SSOs.

1 Introduction

This paper studies the economic and technological impact of voluntary collaborative non-market standard setting organizations (SSOs). SSOs are an important catalyst for coordination in many industries where consumers value inter-operability (e.g. telecommunications and computing). They provide a forum for collective decision-making and an alternative to standardization through market competition or government regulation. SSOs are a diverse set of institutions, ranging from large industry associations to small consortia, and they are involved in a variety of activities, including collaborative R&D, compatibility testing, and product certification. While the central goal of SSOs is to promote industry coordination by endorsing a particular technology, these organizations rarely have any formal powers to enforce their recommendations. As a result, SSOs work to create a consensus around particular technologies that can serve as a focal point for industry coordination or lead to a bandwagon process among adopters.

Substantial resources are devoted to the standard setting process as SSOs have grown in both economic importance and policy relevance (Cargill 1997; Shapiro 2000). For example, SSOs played a central role in antitrust actions taken by the U.S. Federal Trade Commission against Dell and Rambus—two frequently cited cases in the debate over US patent reform (Jaffe and Lerner 2004; Farrell et al 2004). Another example is the Standards Development Organization Advancement Act of 2004 (H.R. 1086), which provides greater antitrust immunity to registered SSOs. Much of the policy interest in this area centers on developing a better understanding of how SSOs influence technology choice and market competition.

Our paper is the first to provide a general and systematic measurement of the economic and technological impact of these institutions. Evaluating the role of SSOs is difficult because they operate in diverse markets and their effect on such standard variables as price and quantity is ambiguous. However, a ubiquitous problem for SSOs is the treatment of intellectual property rights (IPR). Participants in the standard setting process are usually obliged to disclose relevant patents to the SSO. In this paper, we use these patents as a window into the role of SSOs in technological innovation. This approach builds on a prior literature that has established the validity of patent citations as a measure of the economic and technological significance of the underlying innovation (Jaffe and Trajtenberg, 2004; Hall, Jaffe and Trajtenberg, 2005).

Our analysis of SSO performance begins with a sample of 863 IPR disclosures from four major SSOs: the American National Standards Institute (ANSI), the Institute for Electrical and Electronic Engineers (IEEE), the Internet Engineering Task Force (IETF), and the International Telecommunications Union (ITU). These disclosures referenced a total of 657 U.S. patents, which we merged with the NBER U.S. patents database (Hall, Jaffe and Trajtenberg 2001).

Our first look at citation patterns reveals that SSO patents receive many more citations than an average patent from the same technological field and application year. Not surprisingly, SSO patents are more important than the average patent. A more striking result uses methods developed by Mehta, Rysman, and Simcoe (2006) to demonstrate a significant difference in the age distribution of these citations. Specifically, SSO patent citations are less concentrated in the first few years after the patent is granted—suggesting that these patents are both more significant and have a longer useful life than the average patent.

Why do the SSO patents exhibit a different citation-age distribution? We consider two possible explanations—SSOs might select patents corresponding to important technologies, or they may cause patents to exhibit the citation profile we observe. The selection effect is natural given that SSOs explicitly attempt to identify the best technology to serve a given need. The causal effect may arise because an SSO embeds a technology in a standard that exhibits long-lasting economic importance because of network effects and path-dependence, or because an SSO disclosure represents a public announcement that attracts attention to a patent and creates bandwagons in the technology adoption process.

Distinguishing between the selection and causal effects requires the estimation of a counterfactual: what would have happened to a disclosed patent if the disclosure had never occurred? We consider two approaches to this problem. The first approach focuses on SSO patents and uses pre-disclosure observations to estimate the counterfactual citation rate. In this model, the impact of disclosure is identified by within-patent changes in citation frequency following disclosure. Our second approach combines the SSO patents with a set of "controls" in a pooled cross-sectional regression. This allows estimation of both a time-invariant SSO effect (selection) and a time-varying disclosure effect (coordination benefits).[1] Both methods rely on variation in the age of patents when they are disclosed. While we cannot sign the potential bias from measurement error or endogeneity of the disclosure date, the main results do not change when we vary our assumptions about the timing of disclosure.

We find that the baseline citation rate for SSO patents is roughly three times that of the average patent. We also find that disclosure produces a 30 percent increase in the SSO patent citation rate. These results indicate that SSOs select technologies that are already important *and* increase their significance through formal endorsement and other efforts to promote industry coordination. Although it is difficult to attach a dollar value to citation counts, the estimates in Harhoff et al (1999) and Hall, Jaffe and Trajtenberg (2005) suggest that they are economically meaningful. By extension, our findings suggest that SSO endorsements

[1] In our regressions, the "selection effect" measures differences between an average patent and an SSO patent. This could be larger or smaller than the difference between a patent "at risk" for disclosure (i.e. a patent on a technology that is evaluated by an SSO), and a patent that is essential to implement the formal standard.

can increase the value of patented technology—a central claim in recent policy debates over rent-seeking and the governance of the standards creation process.

This paper adds to a small body of economic research on SSOs, which includes theoretical contributions by Farrell and Saloner (1988), Lerner and Tirole (2005) and Farrell and Simcoe (2006), and empirical studies by Chiao, Lerner and Tirole (2005), Simcoe (2005) and Gandal, Gantman and Genesove (2005). Our findings also contribute to a growing empirical literature that examines the impact of particular institutions on the process of technological change. Examples of this research include Furman and Stern's (2004) study of biological resource centers, and studies of the university-industry interface, including Henderson, Jaffe and Trajtenberg (1993) and Mowery et al (1999).

In the next section, we describe the four SSOs examined in this paper and how they treat intellectual property. Section 3 describes the data set, while Section 4 takes an initial look at the difference in citation patterns between the SSO and control samples. Section 5 examines the post-disclosure increase in citation rates. Section 6 offers some conclusions.

2 SSOs and Intellectual Property

This paper examines patents disclosed in the standard-setting process at four major SSOs. These groups are the American National Standards Institute (ANSI), the Institute of Electrical and Electronics Engineers (IEEE), the Internet Engineering Task Force (IETF), and the International Telecommunications Union - Telecommunication Standardization Sector (ITU-T, or often, ITU). The ITU is an international institution focused primarily on telecommunications standards. While international in scope, the IEEE and IETF draw the majority of their participants from North America, and are usually associated with the computer hardware and software industries (although some of their most significant standards are communications protocols). ANSI is an umbrella organization that promulgates a common set of rules and procedures for U.S. standards developers in a wide variety of different industries. However, the majority of IPR disclosed to ANSI certified SSOs also comes from the computing and communications industries. Table 1 provides some indication of the scope of these organizations based on the technology class assigned to each patent in our data.

The ITU is the oldest of our four SSOs, with origins dating back to around 1865. Its original mission was to promote international coordination among the various rapidly expanding domestic telephone networks. The ITU is based in Switzerland, and its membership consists of delegates from member nations along with representatives of the larger firms or network operators in each of these countries. The ITU's standard setting activities continue to emphasize the protocols used to operate the international telephone network. Recent efforts have

Table 1: Technology Classification of SSO Patents†

	ANSI	IEEE	IETF	ITU	Totals
Computers & Communications	40	113	29	71	253
Computer Hardware & Software	61	91	58	72	282
Computer Peripherals	4	0	1	0	5
Information Storage	10	7	2	0	19
Electrical Devices	1	7	0	2	10
Electrical Measure & Test	4	2	0	1	7
Semiconductor Devices	0	9	0	0	9
Misc. Electrical	4	1	0	40	45
Material Processing	8	0	0	0	8
Optics	8	1	0	9	18
Others	12	5	3	5	25
All Categories	152	236	93	200	657
	Overlap in Patent Disclosures				
ANSI overlap	152	5	7	16	
IEEE overlap	5	236	10	0	
IETF overlap	7	10	93	5	
ITU overlap	16	0	5	200	

†Based on subcategory classifications in the NBER U.S. patent database.

focused on numbering and addressing, network services, physical interconnection, monitoring and accounting, traffic management, and quality of service.

The IEEE is slightly younger than the ITU. It was founded in 1884 by several pioneers in the field of electrical engineering. Although the IEEE is a professional society whose members are individual engineers, it is possible to become a corporate member when participating in its standard setting activities. The IEEE's standard setting efforts cover a wide range of subjects, from electrical safety, to cryptography, to standards for semiconductor testing equipment. In recent years, the IEEE's most commercially significant standards work has revolved around the 802.11 specifications for wireless computer networking.

ANSI was formed in 1918 to coordinate the ongoing standards development efforts of a number of different organizations.[2] ANSI continues to play a role in coordinating the activities of hundreds of different U.S. SSOs—primarily through an accreditation program focused on key dimensions of the standards development process.[3] While the IEEE is an ANSI certified

[2]The original ANSI members were the American Institute of Electrical Engineers (now IEEE), the American Society of Mechanical Engineers (ASME), American Society of Civil Engineers (ASCE), American Institute of Mining and Metallurgical Engineers (AIMME), and the American Society for Testing Materials (ASTM).

[3]ANSI also serves as the U.S. representative on the two major non-treaty international standards organizations, the International Organization for Standardization (ISO) and the International Electrotechnical Commis-

SSO, Table 1 shows that the majority of the patents that we identified in ANSI's disclosure records came from other organizations. In fact, many of the ANSI disclosures are associated with the Telecommunications Industry Association, which has worked on technologies such as DSL (for data transmission over phone lines) and TDMA (a cellular telephony protocol). [4]

The IETF is the least formal of the four SSOs studied in this paper. This organization grew out the ARPANET engineering community that emerged during the 1970s, and did not resemble a formal SSO until the late 1980s or early 1990s (Mowery and Simcoe, 2002). The IETF creates a host of protocols used to run the internet. Prominent examples include the Internet's core transport protocols (TCP/IP and Ethernet), standards used to allocate network addresses (DHCP), and specifications used by popular applications such as e-mail or file transfer. From its inception, membership in the IETF and its various working groups has been open to any interested individual. Much of the IETF's work takes place in online forums sponsored by individual committees and is visible to the general public.

While these four SSOs differ in terms of their technology focus, membership rules, and level of formality, the broad outlines of their standard setting process are quite similar. This process always begins with the recognition of some coordination problem, which leads to the formation of a technical committee. The committee's job is to analyze the problem and recommend a consensus solution. The process of identifying and evaluating alternative solutions often lasts for several years. When the committee reaches a consensus, the SSO will issue a formal endorsement—which hopefully serves as a catalyst for widespread implementation and adoption of a new industry standard. Some SSOs also encourage diffusion through marketing and certification activities.

Intellectual property rights are an increasingly important part of the technology evaluation process at many SSOs. While most SSOs would like to avoid the distributional conflicts and obstacles to implementation that patents can produce, IPRs are increasingly an unavoidable part of technology development and commercialization. In part, this reflects the well-documented surge in patenting that began in the mid-1980s. However, it also reflects the fact that many firms would like to own IPR that is embedded in an industry standard. Patent owners frequently seek royalty payments for the use of their technology—even (or, perhaps, especially) when it is essential to the implementation of an industry standard. This creates strong incentives to push for one's own technology within the SSO.

SSOs' intellectual property polices consist of three basic pieces—search, disclosure, and

sion (IEC).

[4]ANSI does not require approved SSO's to forward disclosures to ANSI and so the ANSI sample represents just disclosures that a member SSO chose to forward. That explains why there is little overlap in Table 1, even though the IEEE is a member of ANSI. While this feature changes the interpretation of the ANSI sample, it is useful that it looks to largely independent sets of patents.

licensing rules. Lemley (2002) presents a survey of IPR polices from a large number of SSOs. All four of the SSOs examined in this paper use variations on the relatively common policy of "reasonable and non-discriminatory" licensing (RAND). Under this policy, SSO members agree to disclose any known property rights as soon as possible. They are not, however, obliged to carry out a search. When a patent or other piece of intellectual property is discovered, the SSO seeks assurances from the owner that they will license the technology to any interested standards implementor on reasonable and non-discriminatory terms.[5] While SSOs and their individual technical committees are generally inclined to search for technologies that are unprotected or available on a royalty-free basis, their job is to evaluate the potential trade-off between technical quality and openness.

IPR disclosures often occur near the end of the standard setting process, when the SSO is close to endorsing a particular technology. In some cases, firms withhold their IPR for strategic reasons. For example, they may have an unpublished pending patent application, or believe that disclosure would lead potential implementors to reconsider technology-specific investments that would increase the value of the patent. There are also practical reasons for firms to delay their patent disclosures. In particular, it is often possible to save money by delaying a full patent search until the outlines of a final specification become clear.

Figure 1 illustrates the growth in intellectual property disclosure at the four SSOs that we study. (We define a disclosure as an announcement by a single firm on a given date that it potentially owns one or more pieces of intellectual property needed to implement a proposed standard.[6]) While the number of IPR disclosures was initially quite small, it began to grow during the early 1990's. By the late 1990s, all four SSOs were experiencing significant growth.

For our purposes, the rise in IPR disclosure means that we have access to a publicly available list of patents associated with specific SSOs. Many features of these patents—such as the number of citations they receive—are easily compared across different industries and time periods. Thus, they provide a unique window through which to examine the economic and technological significance of SSOs.

[5]In practice, the "reasonable and non-discriminatory" requirement in a RAND licensing policy seems to imply very few obligations on the part of prospective licensors. The reasonableness requirement is rarely taken to mean that the technology must be offered at a uniform price. When the intellectual-property holder has not made an *ex ante* commitment to some set of licensing terms, each potential implementor of the standard will negotiate their own terms. And while licensors are expected to negotiate in good faith with any potential developer, the individual terms offered may vary widely. SSOs have been very hesitant to get further involved in the negotiating process. In part this reflects their own concerns about the antitrust implications associated with any type of collective pricing agreement. At the same time, it also likely reflects their fear of alienating particular members.

[6]When a firm claims that a single patent covers two or more standards, each one counts as a separate disclosure. However, we only keep one copy of the patent in our data for analysis.

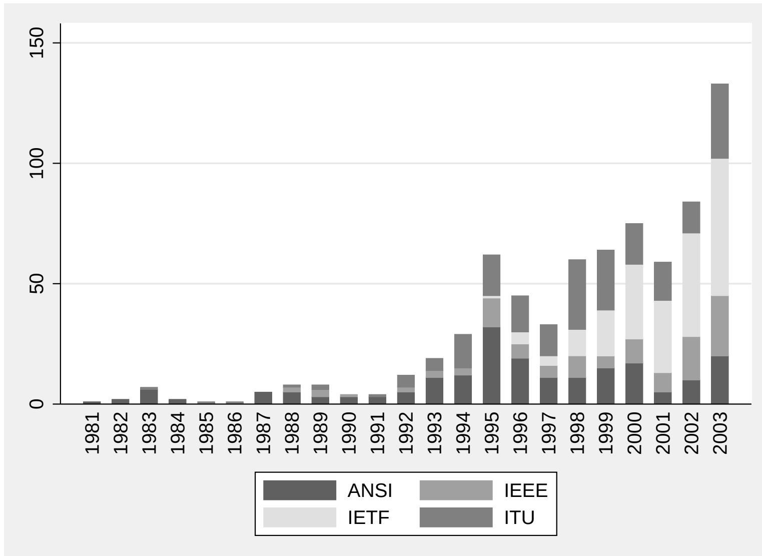

Figure 1: Intellectual Property Disclosures

3 Data

At most SSOs, an IPR disclosure consists of a letter or email to the SSO, indicating that some company owns (or may own) intellectual property that could be relevant to a proposed standard. We identified 863 disclosure letters within the publicly available records of the four SSOs in this study. While these disclosures begin in 1971 and continue through 2003, Figure 1 shows that the majority occurred during the 1990s.

A close examination of the disclosure letters reveals that their contents often vary dramatically—both within and between SSOs. Some disclosures contain detailed licensing terms and refer to specific patents, while others are simply general statements regarding a firm's willingness to offer a RAND license should they own any relevant intellectual property. Overall, this variation in practice reflects differences in SSO participants, policies, and objectives, as well as evolving industry norms with respect to the entire issue of disclosure. Table 2 presents several summary statistics that illustrate the amount of variation within our disclosure sample.

Though the IETF is the last group in our sample to receive a formal IPR disclosure, Table 2 shows that they received more disclosures than any other SSO. However, disclosures to the IETF contained the smallest number of references to a U.S. patent number. This puzzling result reflects differences in the size of disclosures (i.e. the amount of IPR disclosed in a given letter), which are driven by three factors—the number of multi-national "patent families" disclosed, the scope of the proposed standard, and the specificity of individual disclosures. Specificity

Table 2: IPR Disclosure Summary Statistics

	IPR Disclosure Summary				Patent Counts	
	First Disclosure	Total Disclosures	Average Size†	Average Specificity††	U.S. Patents	Total Patents
ANSI	May, 1971	223	2.88	0.39	184	211
IEEE	January, 1988	121	3.55	0.83	315	387
IETF	June, 1995	295	1.53	0.23	123	138
ITU	October, 1983	224	3.68	0.58	433	562

†Size is a count of the patents/applications mentioned in the disclosure.
††Specificity equals one if the disclosure provides one or more patent/application numbers that uniquely identifies the relevant IPR.

refers to the probability of listing a particular patent or application number that can be used to identify the relevant IPR. (For example, it was a common practice at the IETF for several years to "disclose" the existence of an unpublished patent application without providing any information that could be used to verify its existence.) Table 2 indicates that ANSI, IEEE and the ITU have a greater average disclosure size and specificity than the IETF.

The last two columns in Table 2 show the number of patents referenced in the disclosure letters. While the majority of these patents were issued in the U.S. a number of international patents were disclosed at each SSO. These patents are often part of an international "family" whose U.S. counterpart appears in our estimation sample. Many IPR disclosures refer to unpublished patent applications, which we were not able to link to the NBER data. Table 1 also shows a small amount of overlap created by patents disclosed to more than one SSO. After removing all of the foreign patents, patent applications, and duplicate observations, the intellectual property disclosures made at ANSI, the IEEE, IETF and ITU yield a pooled sample of 657 unique U.S. patents.

Before turning to these patents, it is important to note that the disclosure data have several limitations. First, while it is trivial to link an IPR disclosure to an SSO, it is often quite difficult to make the link to a particular standard. As a result, we observe only disclosures—not whether the proposal became a standard, or whether the IPR was essential to the final specification (i.e. whether an implementor of the standard would need to license the disclosed patent). And second, it is unlikely that disclosed patents are broadly representative of the technology evaluated by these four SSOs. Rather, they are likely to be among the most important patents owned by the disclosing firm, and to be concentrated within several of the most commercially significant standard setting efforts. Nevertheless, we believe that these patents provide a unique window into the technology evaluated by SSOs, and can be used to address important questions about SSO performance.

We begin our evaluation of the SSO patents by linking them to the NBER U.S. patent data file (Hall, Jaffe, Trajtenberg 2001), which contains several important variables, including application and grant dates, assignee names, and citation counts.[7] Table 3 compares the SSO patents to a set of control patents with a matching application year and primary 3-digit USPTO technology classification. This comparison shows that SSO patents contain more claims and are more likely to be part of a "family" of patent applications spanning multiple countries. Prior research has shown that these variables are positively correlated with a patent's economic value. The table also shows that SSO patents are more likely to be assigned to a U.S. company than the control patents.

Table 3: SSO Patent Characteristics

	Pooled Sample		Individual SSOs			
	SSO	Controls†	ANSI	IEEE	IETF	ITU
Total Claims	21.0 (17.9)	14.7 (12.4)	20.4 (15.8)	23.5 (21.8)	22.83 (17.9)	17.8 (12.7)
International Family	0.40 (0.49)	0.30 (0.46)	0.30 (0.46)	0.40 (0.49)	0.27 (0.45)	0.53 (0.50)
Assignee Type						
US Company	0.73	0.50	0.78	0.73	0.72	0.65
Foreign Company	0.20	0.40	0.14	0.18	0.14	0.31
Other	0.07	0.10	0.08	0.09	0.14	0.04
Application Year	1991.4 (7.1)	1993.1 (6.5)	1987.3 (9.5)	1993.1 (5.3)	1994.1 (4.3)	1991.1 (6.2)
Patents	633	140,776	152	240	93	173

†The control sample contains all patents that match the application-year and primary 3-digit USPTO technology classification (nclass) of any SSO patent.

While the control patents in Table 3 serve as a useful point of reference, it is unlikely that they are a valid set of "controls" in the sense that they are statistically indistinguishable from a pre-disclosure SSO patent. Our analysis uses the control patents to address macro changes to the patenting regime, and our main results are based largely on variation within the SSO sample. When we compare SSO patents to the control sample, it will be with an eye towards comparing SSO patents to "average" patents, rather than patents that are truly identical but for disclosure.

In the remainder of the paper, our primary measure of economic and technological sig-

[7]The NBER data have been updated through 2002 and are available on Bronwyn Hall's web site http://emlab.berkeley.edu/users/bhhall/bhdata.html. We are also grateful to Ajay Agrawal for providing us with data on citations from patents granted in 2003 and 2004.

nificance is based on forward-citations (i.e. the citations received by a particular patent). A number of papers indicate that forward-citations are a valid measure of economic and technological significance. For example, Hall, Jaffe and Trajtenberg (2005) show that citation weighted patent counts are more correlated with a firm's market value than un-weighted patent counts. Other papers, such as Jaffe, Henderson and Trajtenberg (1993) interpret these citations as an indicator of knowledge transfers from the cited to the citing patent.[8] For this paper, it is not important to defend any particular interpretation of the meaning of a patent citation. As long as forward citation counts contain some information about the technological or commercial significance of the cited invention, we can use them to learn about the impact of SSOs.

While we would like to study the long run impacts of SSO affiliation, we limit the analysis to a period of about 15 years due to data availability. In particular, we have very few observations on "old" SSO patents because the majority of them were either granted or disclosed near the end of our sample period. For example, only 470 of the 657 SSO patents in our SSO sample were disclosed prior to 2002—our final year of citations data. Figure 2 shows the application-year distribution for SSO patents. Most of these patents were not granted until the mid-1990s—reflecting the start of an ongoing surge in both patenting and disclosure. Table 7 (which can be found at the back of the paper) provides counts of the number of patent-year observations in our citations data, both before and after disclosure.

The recent surge in IPR disclosures raises the issue of truncation. The main source of truncation in our data is the lag between the application and grant dates of a citing patent (the grant date is when the citation is observed). Following Hall, Jaffe, and Trajtenberg (2001), we choose to date citations based on the application year of the citing patent. The truncation problems arise because we do not observe those citations from a given application corresponding to patents with long lags. We deal with this issue in two ways. First, we limit our analysis to application years through 2002—even though we collected data for patents granted through 2004—ensuring that we only lose citations associated with a grant lag of three or more years. Second, we include a set of citing-year dummies in all of our regressions.

4 Citation Age Profiles

In this section, we examine the distribution of forward-citations to patents in the SSO and control samples, focusing on the citation age profile—i.e. the average citation rate conditional

[8] This interpretation raises the question of how to treat self-citations (i.e. citations to a patent owned by the same entity as the citing patent). We found that there was little difference in the results presented below when self-cites were excluded from the sample.

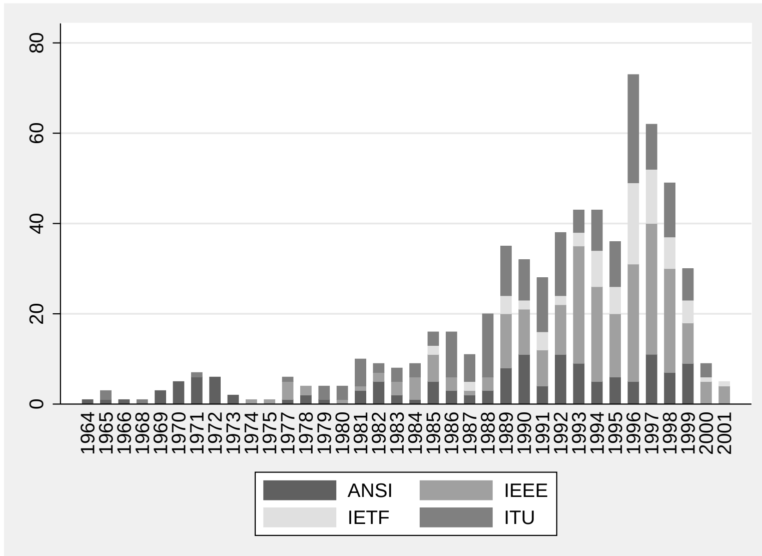

Figure 2: SSO Patent Application Year Distribution

on the age of the cited patent.[9] We begin with a direct comparison of the average citation rates for SSO and control patents before turning to an econometric model that includes application- and citing-year fixed effects to control for a number of confounding factors.

Figure 3 illustrates this section's two main results. First, SSO patents are cited far more frequently than controls.[10] This difference in citation rates is both substantial and persistent. Second, the shape of the citation age-profile is different for the SSO patents. In particular, the peak citation age for SSO patents is later, and the SSO patents receive a larger share of their cumulative citations in later years.

We find these citations patterns interesting for several reasons. The large difference in average citation rates suggests that the technology disclosed to SSOs is quite significant. The market value regressions in Hall Jaffe, and Trajtenberg (2005) also indicate that the "unexpected future citations" reflected in a flatter SSO age-profile are more valuable than an average citation. Finally, the fact that citations to SSO patents differ from control patents suggests two competing hypotheses: either SSOs cause an increase in the citation rate, or they select patents on the basis of an expected increase in future citations. However, before turning to the question of selection versus marginal effects, we develop an econometric model to illustrate the

[9]Hall, Jaffe, and Trajtenberg (2001) refer to this statistic as the lag distribution.

[10]The SSO patent with the most cumulative citations is number 4,405,829, which covers essential methods for public-key cryptography. Granted in 1983, this patent had received 368 citations by 2002. The inventors on the patent are Ronald Rivest, Adi Shamir and Leonard Adelman—whose initials (RSA) are very well-known in crypotography circles.

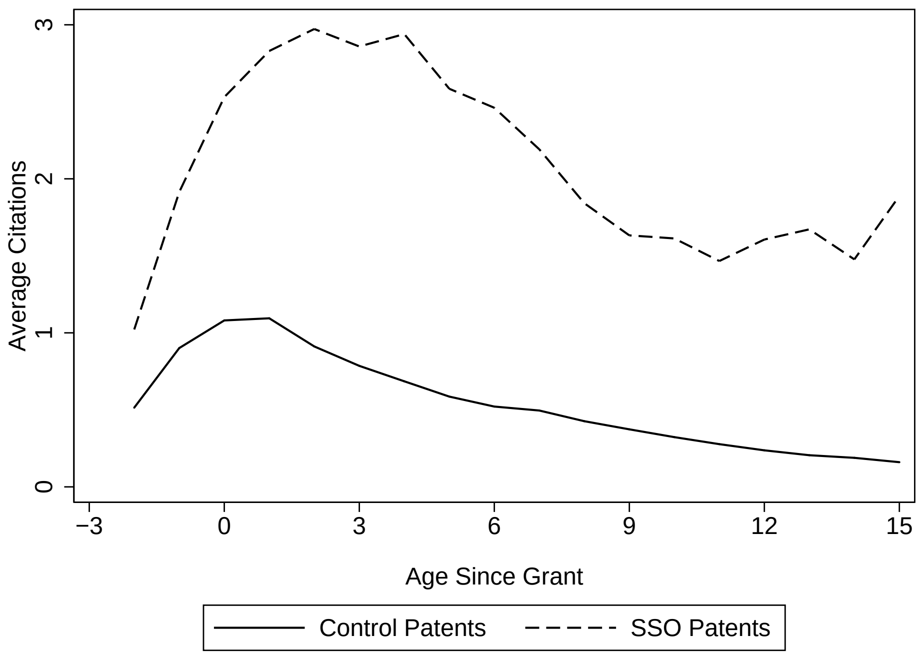

Figure 3: E[Cites|Age] for SSO and Control Patents

substantial difference in the age profile of the SSO and control patents.

We estimate the citation age profile following an approach proposed in Mehta, Rysman and Simcoe (2006). This method uses a full set of application- and citing-year effects to control for various confounding factors—such as policy changes and funding issues at the USPTO, increases in citation propensity over time, and differences in the technological significance or "fertility" of various application-year cohorts. It is well known that one cannot identify a full set of patent-age, application-year and citing-year effects in a linear model (because age equals citation year minus application year). Prior research on the age-profile of patent citations has relied on non-linear functional form restrictions to solve this problem. Mehta, Rysman and Simcoe suggest an alternative approach based on the assumption that the citation age process actually begins when a patent is granted (rather than its application-year). This assumption is reasonable because the age process is meant to capture a process of diffusion and obsolescence. Plausibly, that process does not begin until the information in a patent is publicly available, which is the grant date for U.S. patents. Mehta, Rysman and Simcoe present evidence in support of this assumption. If the publication lag is exogenous, this re-definition of "age" allows for non-parametric identification of the citation age profile. Intuitively, the age effects are identified by comparing the citation rate of patents from the same application-year cohort whose "age" differs as a result of variation in the length of the USPTO review process.[11]

We estimate a set of citation age profiles using the following model, where C_{it} is the number of citations received by patent i in year t, α_y are fixed effects for application year y, α_t are fixed effects for citing year t (as measured by the application year of the citing patent), α_c are fixed effects for the three-digit USPTO technology classification, α_a^{CTRL} and α_a^{SSO} are the age effects for the control patents and SSO patents at age a, ε_{it} is a patent-year error term that is uncorrelated with the fixed effects, and $f()$ is a Poisson process. Here, age is defined relative to the grant year g, i.e. $a = t - g$.

$$C_{it} = f(\alpha_y, \alpha_t, \alpha_c, \alpha_a^{CTRL}, \alpha_a^{SSO}, \varepsilon_{it}) \tag{1}$$

This specification is based on the assumption that the application-year and citing-year effects are identical for the SSO and control sample, but the age profiles can be different. While both the control sample and the SSO sample contribute to identifying the application-year and citing-year effects, the number of observations in the control sample dwarfs the number in the SSO sample. Conceptually, we are using the control sample to identify the application-year and citing-year effects, while estimating a separate age profile for each sample. Hence, the choice of the control sample has little effect of the shape of the SSO age profile.

We estimate Equation (1) separately on the pooled sample and for each SSO. Table 8, which can be found at the end of the paper, provides a complete set of age coefficients from each of these regressions.[12] The most obvious result of this exercise is that the SSO age effects are larger than the controls. This is not surprising given that most of the control patents receive very few citations (as can be seen in Figure 3). Still, the absolute difference in citation rates is striking.

Since it is difficult to evaluate hypotheses about the shape of the age distribution using the coefficients in Equation (1), we rely on summary statistics. In particular, we predict the number of citations conditional on age (setting the dummy variables for application year 1999 and citation year 1999 on and leaving all other application and citation years off) and use these values to compute a probability distribution. Then, we use the probability distribution to compute an "average citation age" for each group of patents. We compute standard errors

[11] When "age" is defined relative to the grant-year of a patent, it is natural for some patents to receive citations at negative ages. This occurs whenever the application-year of the citing patent is less than the grant-year of the cited patent. For the assumption that age begins at grant date to be exactly correct, it must be that these citations are added by the patent examiner or turned up in a patent search as opposed to indicating an actual intellectual debt. Mehta, Rysman and Simcoe (2006) discuss this at length. In practice, we drop citations from ages below -2 from our data set.

[12] One patent disclosed to the IETF has an application year of 1977 while all the rest are applied for in 1985 or later. We drop the 1977 patent in the following analysis.

for this statistic using the delta method, and test the hypothesis that the mean citation-age is equal in the SSO and control samples.[13]

Figure 4 graphs the citation probability distributions over ages -2 to 12 as computed from the regression results. In each case, we can see that the SSO distribution is lower at low ages and higher at high ages. The IETF exhibits the most remarkably long-lived citation profile. Hall, Jaffe, and Trajtenberg (2001) draw similar graphs for a number of groups of patents and always find peaks in the 4th or 5th year after application. This is consistent with our control groups, which show peaks 1 to 2 years after the grant year. However, it contrasts with the SSO patents—particularly the IEEE and IETF—whose peak citation rates occur later.

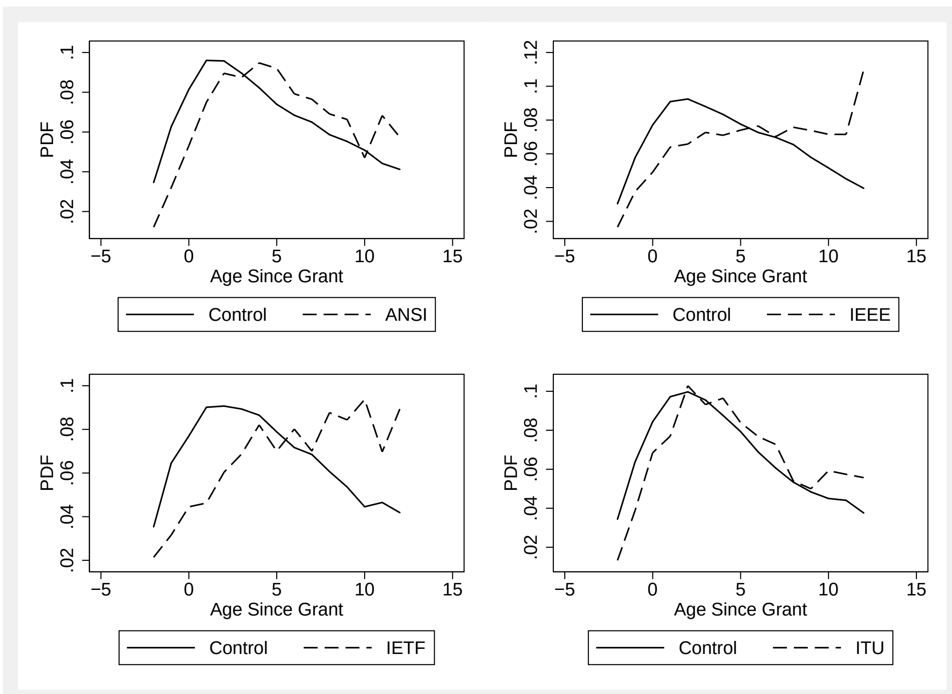

Figure 4: Age profile of citations based on regression results

Table 4 presents estimates of the "average citation age" using both the unadjusted age distribution and the regression model described above. The average age is naturally higher when we use the regression procedure, since it corrects for the truncation problem inherent in observing many patents near the end of the sample period. The important point is that both methods show that SSO patents receive significantly more of their cumulative citations in later years.

[13]We use a heteroskedasticity-consistent variance-covariance matrix (clustered on patent) to perform these calculations.

Table 4: Average Age (since Grant) of Patent Citations

	Raw Data		Estimated PDF		
	Control	SSO	Control	SSO	Difference
Pooled Sample	2.42 (0.00)	4.45 (0.04)	4.49 (0.08)	5.51 (0.21)	1.02 (0.20)
ANSI	2.40 (0.01)	5.88 (0.08)	4.44 (0.16)	5.37 (0.35)	0.87 (0.31)
IEEE	2.01 (0.01)	5.07 (0.07)	4.60 (0.14)	6.04 (0.37)	1.36 (0.37)
IETF	0.94 (0.01)	3.30 (0.06)	4.48 (0.18)	6.19 (0.34)	1.81 (0.31)
ITU	2.43 (0.01)	4.25 (0.05)	4.28 (0.13)	5.04 (0.30)	0.76 (0.27)

The "Estimated PDF" is based on fitted values from the Poisson QML regression model of Equation (1). Standard errors for the average age and difference were calculated using the delta method.

One concern with the estimates in Table 4 may be that the high average citation age in the SSO sample simply reflects greater overall importance. In other words, all highly cited patents might have a similar age profile. In fact, the opposite is true. When we compared the SSO patent age profiles to a set of highly cited controls we found that the difference in age actually increased slightly. [14] We believe that the explanation for this result is that the plurality of patents get no citations, which implies a flat age profile. It is the patents that actually get citations that generate the hump-shaped age profile. Removing the patents that get no citations from the control sample simply exaggerates this shape.

5 The Impact of SSOs

The previous section showed that patents disclosed to SSOs are cited more often than an average patent and at later ages. Both of these findings suggest that the SSO patents embody significant inventions. However, these results have two plausible interpretations. Differences between the SSO and control patents could simply be a selection effect, whereby SSOs identify and endorse technologies that are more likely to exhibit a particular age-profile. On the other hand, differences in the citation age profile may reflect the causal impact of an SSO endorsement

[14]We defined highly cited patents to be those that were in the top 10% of citations received over the life of the patent relative to other patents in the same technology class and application year. This cut-off created a control sample with an average citation rate slightly higher than the SSO sample.

on the significance of the underlying technology. In this section, we address this question by studying the relationship between citation rates and the timing of disclosure. Our goal is to estimate the marginal impact of the SSOs on the patent citation rate.

We use two different methods to estimate the disclosure effect. Our first approach uses only those patents disclosed to an SSO, and relies on variation in the timing of patent disclosures to identify the marginal effect. For this experiment, we do not use the control patents. The second approach relies on a pooled cross-sectional specification similar to the age-profile regressions presented above. However, we include an SSO dummy to estimate the selection effect and a post-disclosure dummy to estimate the marginal impact of the SSO. In order to estimate a single SSO dummy, we restrict the age process to be the same for the SSO and control samples. Although this is a strong assumption, doing so allows us to make a compelling comparison between the selection and marginal effects. Remarkably, the two approaches produce very similar estimates of the marginal impact of an SSO endorsement—disclosure is associated with a 20 to 40 percent increase in the citation rate.

5.1 Marginal Effects in the SSO Sample

In this sub-section, we use variation in the timing of SSO patent disclosures to estimate the marginal effect. Specifically, we ignore the control patents and use pre-disclosure SSO patents to estimate a counterfactual citation rate for disclosed patents. Since we are no longer interested in separating the age, cohort and calendar effects, we rely on a more flexible specification that includes individual patent fixed-effects. Specifically, we estimate the following fixed-effects Poisson model, where α_{it}^{Disc} is a post-disclosure dummy that captures the marginal effect of the SSO; α_a are a set of patent-age effects; α_{trunc} are a set of year dummies for 1998 through 2003, which control for truncation bias induced by the lag between application and grant; and γ_i is a patent conditional fixed effect.[15]

$$C_{it} = f(\alpha_{it}^{Disc}, \alpha_a, \alpha_{trunc}, \gamma_i, \varepsilon_{it}) \tag{2}$$

By removing the control patents and introducing patent-level fixed effects, this regression addresses any concerns about the selection of SSO patents based on time-invariant unobserved characteristics. In particular, α_{it}^{Disc} is estimated entirely off of within-patent variation in citation rates and between-patent variation in the timing of disclosure. For example, if all SSO

[15] Wooldridge (1999) shows that this estimator is consistent under quite general conditions (i.e. only a conditional mean assumption is required) and robust to arbitrary serial correlation in the dependent variable. Stata code for computing the robust standard errors is available at http://www.rotman.utoronto.ca/timothy.simcoe/xtpqml.txt

patents were disclosed at the same age, α_{it}^{Disc} would not be identified since it would be co-linear with the age effects.

The patent fixed-effects and age effects in Equation (2) are collinear with cohort and calendar effects, so we cannot interpret the age effects as such (regardless of how age is defined). However, the patent fixed-effects are more flexible, which is useful since our goal is to focus attention on the post-disclosure coefficient. While it is not possible to estimate a full set of citing-year effects because of collinearity, we include a set of citing-year dummies α_{trunc} for the last five years in our sample to control for truncation (i.e. unobserved citations) introduced by the patent application process. We chose these dates by examining the empirical grant-lag distribution and observing that 99 percent of all observed patents are granted within 5 years of application.

Our interpretation of the post-disclosure parameter as an estimate of the causal impact of the SSO on citation rates rests on the assumption that disclosure timing is exogenous. If disclosure timing is not exogenous, the sign of the associated bias is difficult to predict. For example, suppose there is a large causal effect of disclosure but either SSO participants or firms in the technology market can predict which patents will be disclosed. In that case, patents may begin to receive citations before disclosure, which would cause the correlation between disclosure and citations to understate the impact of the SSO. On the other hand, patent disclosures may be correlated with time-varying unobservables. If SSOs can accurately forecast an increase in citations using information that is not available to us—and if they use these forecasts in selecting a technology to endorse—we will observe an increase in citations around the date of disclosure even if the SSO has no "true" marginal impact.

It is not possible to test the assumption that disclosure dates are exogenous. However, we can look for evidence of a pre-disclosure increase in citations. Our baseline model uses a simple post-disclosure dummy to estimate α_{it}^{Disc}. The advantage of using the disclosure year as our break point is that disclosures tend to occur near the end of the standard setting process—when it has become relatively clear which technologies will be used in the standard.[16] This is a logical place to begin looking for a network or bandwagon effect.

However, we also consider what happens if the post-disclosure dummy is activated two years before the actual IPR disclosure. There are several reasons why an SSOs' impact on patent citations might coincide with the start of the standard setting process rather than the actual disclosure event. Firms may be able to anticipate the SSO's technology choice. There may be

[16]Firms have both strategic and practical reasons to delay disclosure until the end of the standard setting process. Firms hoping to use the standard-setting process to create hold-up opportunities naturally want to wait until technology-specific investments have been made before revealing their IPR. However, delay also allows firms to undertake a more focused and less costly patent search.

substantial lags between the emergence of a consensus (which starts the technology adoption bandwagon) and the SSO's formal decision (which triggers the intellectual property disclosure). Finally, since we date citations to the application-year of the citing patent, some citations that should follow the disclosure event—based on when they are added to the patent—may appear to precede it.

Table 5 presents our estimates of the disclosure effect. (We do not report any of the age or truncation effects, all of which were jointly significant.) Interpretation of the estimates from Equation (2) is straightforward. The regression coefficients provide a reasonable first-order approximation of the percentage change in the citation rate. For larger coefficients (e.g. above 0.3) the incidence rate ratio, $\exp(\alpha^{Disc}) - 1$, provides a slightly better approximation. Our main results are based on the pooled sample of SSO patents. Given that we are working with relatively small numbers of patents, we feel that the pooled estimates are less sensitive to outliers and timing issues than the individual SSOs.[17] However, we also present results from each of the individual SSOs for comparison.

Table 5: Marginal Effects in the SSO Sample

DV = Cites_{it}	**Pooled Sample**	ANSI	IEEE	IETF	ITU
	Model 1: Baseline				
PostDisclosure	0.174 (0.067)***	0.292 (0.139)**	0.045 (0.084)	0.152 (0.101)	0.263 (0.130)**
Patents	567	121	224	91	184
Observations	6,042	1,462	2,224	831	2,164
	Model 2: Marginal Effect Starts at Disclosure_{-2}				
$\text{PostDisclosure}_{-2}$	0.166 (0.066)***	0.288 (0.169)*	0.107 (0.091)	0.322 (0.114)***	0.353 (0.127)***
Patents	567	121	224	91	184
Observations	6,042	1,462	2,224	831	2,164
	Model 3: Drop 2 year pre-disclosure window				
PostDisclosure	0.310 (0.103)***	0.415 (0.231)*	0.059 (0.118)	0.527 (0.161)***	0.641 (0.202)***
Patents	547	117	214	89	179
Observations	4,884	1,218	1,755	659	1,792

* Significant at 10%; ** Significant at 5%; *** Significant at 1%. Robust standard errors (Wooldridge 1999). Each column is based on a fixed-effect Poisson-QCML regression using the specification in Equation 2. Age, and truncation-year fixed-effects not reported. For pre- and post-disclosure SSO patent sample-sizes, refer to Table 7.

[17] Table 1 showed that there are strong technological similarities across these four organizations.

The first row of Table 5 presents our baseline estimates, which use a standard post-disclosure dummy to estimate the marginal effect. The post-disclosure coefficient for the pooled sample indicates that disclosure is associated with a 19 percent increase in the citation rate. The individual SSO results show a positive and statistically significant disclosure effect at ANSI and ITU—corresponding to an increase of roughly 30 percent in the citation rate. Our estimate of the IETF disclosure effect is positive but statistically insignificant, and the IEEE effect is negligible.[18]

The second and third rows in Table 5 consider models that use alternative definitions of disclosure. In Model 2, we artificially move the disclosure date forward by two years. While this has no impact on the pooled sample results, it leads to an increase in the marginal effect at three of the four individual SSOs (ANSI was essentially unchanged). The largest increase is for the IETF, where the post-disclosure coefficient doubles and becomes statistically significant. These results suggest variation in the amount of measurement error on our post-disclosure variable across the four SSOs in our sample. However, we find the relatively stable pooled sample results reassuring.

Model 3 returns to the standard definition of disclosure, but omits any observations that fall within a 2 year pre-disclosure window. Intuitively, this increases the likelihood that the baseline against which post-disclosure citation increases are measured precedes the start of the standard setting process. Not surprisingly, this also leads to an increase in the estimated marginal effects—in this case for the pooled sample, as well as all four individual SSOs. The increase is quite large in the pooled sample and for ANSI, IETF and the ITU, but remains negligible for the IEEE.

Comparing the results of these three different models suggests that the marginal effect of disclosure on citation rates is somewhere between 18 and 35 percent. While some of the increase clearly predates the actual disclosure letter, the results from Model 1 suggest that this effect continues for several years after disclosure occurs. (We present more evidence on the timing of the disclosure effect below.) While the individual SSO results suggest some variation in the amount of measurement error on our post-disclosure variable, the pooled results are quite robust. Finally, the individual SSO results suggest that our pooled disclosure effect is identified primarily by patents disclosed to ANSI, IETF and ITU.

[18]One potential explanation for the absence of an IEEE or IETF disclosure effect is that these SSOs have a smaller percentage of post-disclosure observations (see Table 7). Also, it is interesting to note that the results divide up along industry lines. That is, ANSI and the ITU deal primarily with the telecommunication industry and the IEEE and the IETF deal primarily with the electronics and computer industries.

5.2 Selection Effects and Pre-Disclosure Cites

The results in the previous sub-section focused on identifying the disclosure effect, which we interpret as the marginal impact of the SSO. However, we might wish to compare the size of the SSO (selection) effect to the size of the disclosure (marginal) effect, or take a closer look at the timing of disclosure. This is not possible when the estimation sample is restricted to SSO patents.

In this sub-section, we pool control and SSO patents in a cross-sectional regression similar to the one used in Section 4. However, we assume that the SSO and control patents have a common set of age effects and include an SSO dummy to estimate the selection effect, along with a post-disclosure dummy to estimate the marginal effect of disclosure. We then replace the single post-disclosure dummy with a series of "age relative to disclosure" effects and examine time-trends in the SSO patent citation rate (relative to the controls) before and after disclosure.

Our results are based on the following specification, where α_y, α_t, α_c, and α_a are application-year, citing-year, technology-class, and age-effects respectively; the parameters of interest are a selection effect α_i^{SSO} and a marginal effect α_{it}^{Disc}; and ε_{it} is a patent-year error term that is uncorrelated with the all of the fixed effects (including selection and disclosure).

$$C_{it} = f(\alpha_i^{SSO}, \alpha_{it}^{Disc}, \alpha_y, \alpha_t, \alpha_c, \alpha_a, \varepsilon_{it}) \tag{3}$$

As above, in order to interpret the disclosure dummy as a marginal effect, the timing of disclosure must be exogenous. However, we naturally interpret the selection of patents to disclose as endogenous. Thus, we do not interpret the SSO dummy to capture the effect of exogenously forcing a patent to be disclosed to an SSO at some time in the future. Rather, we seek to measure the extent to which the endogenous selection process leads to highly cited SSO patents. [19] The other main assumption in this specification is that SSO and control patents have the same pre-disclosure age profile (i.e. that disclosure explains the age-profile results in Section 4). While this is obviously a strong assumption, it allows us to identify the coefficient on an SSO dummy, which we use to measure the selection effect. This allows for a straightforward comparison between the impact of selection and disclosure.

Table 6 presents estimates of the selection and disclosure effects for the pooled sample and each of the four SSOs. We do not report the application-year, citing-year age-since-grant, and technology-class effects—all of which are jointly significant.

[19] Our broad control group corresponds to a broad definition of the selection effect. In reality, "selection" can be thought of in several stages: an SSO recognizes the need for a solution, then considers candidate technologies and then chooses a particular option. While it might be interesting to construct control samples that identify the selection effect relative to intermediate steps in the process, doing so in a convincing way appears challenging and we do not attempt that here.

Table 6: Pooled Cross-sectional Estimates of Selection and Marginal Effects

DV = Cites_{it}	**Pooled Sample**	ANSI	IEEE	IETF	ITU
SSO_i (Selection)	0.910 (0.050)***	0.817 (0.106)***	0.832 (0.081)***	1.277 (0.085)***	0.859 (0.093)***
Disclosure_{it} (Marginal)	0.297 (0.085)***	0.430 (0.155)***	0.341 (0.167)**	0.011 (0.207)	0.266 (0.114)**
Observations	1,044,794	302,242	397,169	174,936	471,662

** Significant at 5%; *** Significant at 1%. Robust standard errors clustered on patent. Each column is based on a Poisson QML regression using the specification in Equation 3. Application-year, citing-year, age, and technology-class fixed-effects not reported. For SSO patent sample-sizes, refer to Table 7.

The pooled sample coefficients indicate that the selection effect is roughly four times as large as the marginal effect, at 148 percent and 35 percent respectively. Thus, our estimates suggest that 20 percent of the difference between SSO and control patents is due to disclosure, while 80 percent is a selection effect. Although we do not have strong priors for this statistic, these estimates strike us as quite reasonable.

Not surprisingly, estimates of the selection effect are positive and precisely estimated for the pooled sample and all four individual SSOs. Conditional on age, technology-class, application- and citing-year, SSO patents receive 148 percent more citations than the average control patent. Within individual SSO's, the selection effect varies from 126 percent (IEEE) to 259 percent (IETF) more citations than the average control patent. This suggests that SSOs are quite good at identifying important technologies. Our estimate of the marginal effect for the pooled sample is positive and significant—indicating that inclusion in the SSO process increases citations by 35 percent. For three out of the four SSOs (ANSI, IEEE and ITU), estimates of the marginal effect are also positive and significant. These estimates range from a 30 percent increase in the citation rate (ITU), to a 54 percent increase (ANSI).

The estimates in Table 6 assume that an SSO's impact on citation rates will begin in the year of disclosure. However, the previous sub-section discussed several reasons why the marginal effect of the SSO might pre-date the formal disclosure of a patent. If this is simply a measurement problem linked to the dating of either disclosures or citations, it will bias our estimates of the true disclosure effect towards zero.

We examine the timing of the increase in citations relative to disclosure by replacing the post-disclosure dummy in Equation (3) with a series of age-relative-to-disclosure effects for the SSO patents, omitting the dummy for one year prior to disclosure. In other words, we estimate a series of "disclosure effects" conditional on the age of the SSO patent relative to its actual disclosure date. (We also drop the SSO dummy since it is co-linear with the full set of age-relative-to-disclosure effects.) This specification allows us to examine the pre- and post-disclosure citation trajectory of the SSO patents relative to the controls. Because this exercise is more demanding on the relatively small sample of SSO patents, we focus on the pooled sample to increase the precision of our estimates.

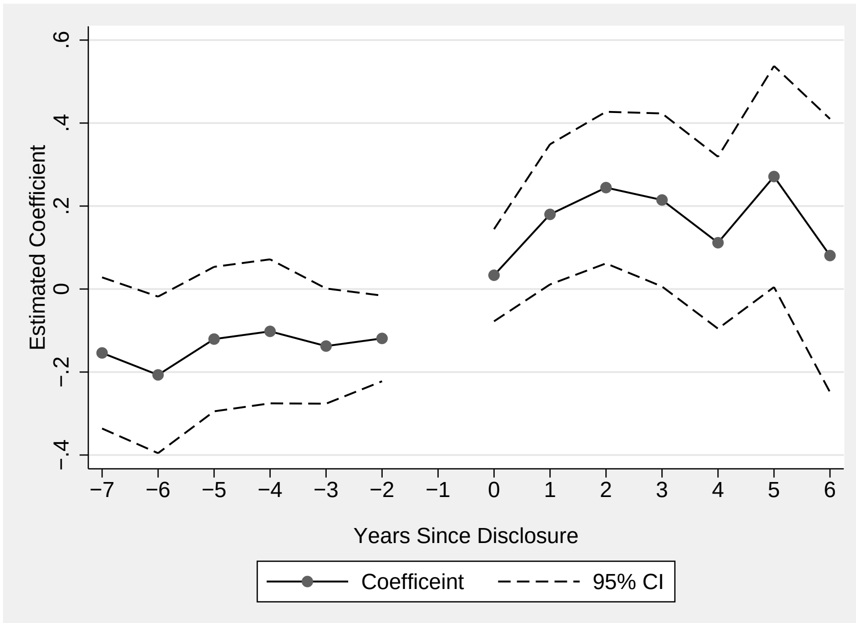

Figure 5: Estimated Pre and Post-disclosure coefficients for SSO Patents

Figure 5 graphs our estimates of the pre- and post-disclosure SSO patent citation-trajectory, along with a 95 percent confidence interval. The relatively flat line from 7 years before disclosure until 2 years before disclosure indicates that the SSO and control patents had similar citation profiles (conditional on age, citing-year, etc.) until about one year before the disclosure event. At this point, the SSO patents experience a fairly sharp increase in citations relative to the controls. After disclosure, the SSO patent citation rate continues to increase relative to the controls for another three years, after which it levels off.

The increase in SSO patent citations during the year before disclosure is roughly 13 percent, while the increase following disclosure is about 27 percent. (The coefficients at -2 and 2 years are -0.12 and 0.24 respectively.) So, the total increase in the SSO patent citation relative to the controls—from two years before disclosure until two years after—is about 40 percent, of which one third pre-dates the actual disclosure.

As we discussed above, there are a number of potential explanations for the observed pre-disclosure "citation bump." In particular, it may provide evidence that SSO patent disclosures are correlated with unobserved measures of the time-varying technological significance of the underlying technology. However, we are encouraged by the relatively flat line from 7 until 2 years before disclosure, which indicates a parallel citation trajectory for the SSO and control patents. To the extent that the timing of disclosures is exogenous, this suggests that the controls actually provide a reasonable estimate of the SSO patents' counterfactual citation rate.

Together, the results in Table 5, Table 6, and Figure 5 suggest that SSOs select important patents *and* have an economically and statistically significant impact on citation rates. In particular, we find that across several different SSOs and estimation methods, citation rates consistently increase by 20 to 40 percent following the disclosure of a patent at the conclusion of the standard setting process. We remain cautious about placing a strong causal interpretation on these results—primarily because it is impossible to test whether firms or SSOs can select patents based on time varying unobserved variables that are correlated with future citations. Nevertheless, lacking any truly exogenous events that push patents into SSO standards, our approach provides a reasonable starting place for identifying the causal impact of SSOs.

Our focus on marginal effects in this section does not imply that we find selection effects uninteresting. Rather, the existence of a significant marginal impact—which we interpret as evidence of bandwagon or network effects—reinforces the importance of identifying and endorsing the best possible technologies. In this respect, the fact that our estimates of the selection effect are relatively large is quite reassuring. Moreover, even if we interpret the marginal effects as evidence of selection on unobservable characteristics, the results in this section would imply that SSOs are remarkably effective at finding "future stars" within a set of technologies that are already highly influential. However one interprets our estimates of the marginal effect, these results show that SSOs play an important role in the process of technological change.

6 Conclusions

The importance of SSOs in technology industries has been widely discussed, and there are many detailed case studies of the formal standard setting process. However, there have been few attempts to systematically measure the impact of these institutions on economic performance or technological change. This paper is the first to address these questions using patent citations as a measure of SSO performance. Our approach leads immediately to the question of causality. In particular, do SSO's influence the process of cumulative technological development, or do they merely identify and evaluate important technologies?

We find substantial evidence that SSOs identify and endorse important technologies. In particular, patents disclosed in the standard setting process receive roughly three times as many citations as a set of controls from the same technology-class and application-year. At the same time, we find a significant increase in the citation rate of SSO patents following disclosure. This marginal effect accounts for roughly 20 percent of the difference in citation rates between SSO and control patents. More importantly, it suggests that SSOs contribute to the ongoing significance of the technologies they endorse through their efforts to promote industry coordination.

Although this paper emphasizes the positive question of SSOs' impact on technological change, our principal findings are relevant to the current policy debates over intellectual property and industry standards. In particular, they suggest that an SSO endorsement has economic value. This implies that we should see firms competing to have their own technologies (and patents) endorsed by these informal groups. While the question of how firms compete for endorsement raises a number of questions that we hope to address in future research, we should acknowledge that it is hard to draw any clear welfare implications from our current results. The impact of having patents in an industry standard will depend on the rules of the SSO, participants' willingness to license any essential intellectual property, and whether they can do so on "reasonable and non-discriminatory" terms.

References

Cargill, Carl F. 1997. *Open Systems Standardization : A Business Approach.* Upper Saddle River, NJ: Prentice Hall PTR.

Chiao, B., Lerner, J., & Tirole, J. 2005. The Rules of Standard Setting Organizations: An Empirical Analysis. *NBER Working Paper #11156.*

David, P. A., & Greenstein, S. 1990. The Economics of Compatibility Standards: An Introduction to Recent Research. *Economics of Innovation and New Technology*, **1**, 3–42.

Farrell, J., & Saloner, G. 1988. Coordination through Committees and Markets. *Rand Journal of Economics*, **19**(2), 235–252.

Farrell, J., & Simcoe, T. 2006. Choosing the Rules for Formal Standardization. *B.E. Press Advances in Economic Analysis and Policy, forthcoming.*

Farrell, J., Pil Choi, J., Edlin, A., Greenstein, S., Hall, B., & Saloner, G. 2004. *Brief Amicus Curiae of Economics Professors and Scholars.*

Furman, J., & Stern, S. 2004. Climbing Atop the Shoulders of Giants: The Impact of Institutions on Cumulative Research. *Working Paper.*

Hall, B., Jaffe, A., & Trajtenberg, M. 2001. The NBER Patent Citations Data File: Lessons, Insights, and Methodological Tools. *NBER Working Paper #8498.*

Hall, B. H., Jaffe, A., & Trajtenberg, M. 2005. Market Value and Patent Citations. *Rand Journal of Economics*, **36**(1), 16–38.

Harhoff, D., Narin, F., Scherer, F. M., & Vopel, K. 1999. Citation Frequency and the Value of Patented Inventions. *Review of Economics and Statistics*, **81**(3), 511–515.

Jaffe, A., & Trajtenberg, M. 2004. *Patents, Citations and Innovations: A Window on the Knowledge Economy.* Cambridge: MIT Press.

Jaffe, A., Trajtenberg, M., & Henderson, R. 1993. Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *Quarterly Journal of Economics*, **108**, 577–598.

Jaffe, Adam B., & Lerner, Joshua. 2004. *Innovation and its Discontents.* Princeton, N.J.: Princeton University Press.

Lemley, M. 2002. Intellectual Property Rights and Standard Setting Organizations. *California Law Review*, **90**, 1889–1981.

Lerner, J., & Tirole, J. 2006. A Model of Forum Shopping with Special Reference to Standard Setting Organizations. *American Economic Review, forthcoming*.

Mehta, A., Rysman, M., & Simcoe, T. 2006. Identifying the Age Profile of Patent Citations. *Social Science Research Network WP#892457*.

Mowery, D., Nelson, R., Sampat, B., & Zeidonis, A. 1999. The Effects of the Bayh-Dole Act on U.S. University Research and Technology Transfer: An Analysis of Data from Columbia University, the University of California, and Stanford University. *In:* Branscomb, L., Florida, R., & Ishii, M. (eds), *Industrializing Knowledge*. Cambridge: MIT Press.

Mowery, D. C., & Simcoe, T. 2002. Is the Internet a US invention? - an economic and technological history of computer networking. *Research Policy*, **31**(8-9), 1369–1387.

Shapiro, C. 2000. Setting Compatibility Standards: Cooperation or Collusion? *Haas School of Business Working Paper*.

Simcoe, T. 2005. Standard Setting Committees. *Social Science Research Network WP#899595*.

Wooldridge, J. M. 1999. Distribution-free estimation of some nonlinear panel data models. *Journal of Econometrics*, **90**(1), 77–97.

Table 7: SSO Patent Observations by Age* (Pre & Post Disclosure)

	Pooled Sample			ANSI		IEEE		IETF		ITU	
	Pre	Post	Disclosed†	Pre	Post	Pre	Post	Pre	Post	Pre	Post
Age -2	496	1	1	110	0	198	0	82	0	156	1
Age -1	617	2	55	148	0	236	0	93	0	195	2
Age 0	567	57	97	135	15	222	14	86	7	178	22
Age 1	436	146	56	100	45	183	29	66	20	132	60
Age 2	353	178	59	73	59	154	37	54	23	111	69
Age 3	278	206	32	59	68	121	49	43	22	89	80
Age 4	234	193	43	47	72	105	43	38	16	73	77
Age 5	174	207	33	35	77	78	51	28	15	58	77
Age 6	134	220	21	30	80	51	64	24	13	49	77
Age 7	109	207	19	26	80	37	58	20	12	45	72
Age 8	83	191	16	23	77	25	48	15	10	38	66
Age 9	63	175	10	17	72	18	42	14	6	30	62
Age 10	50	158	8	13	66	11	42	12	4	27	52
Age 11	38	145	11	11	60	9	38	7	5	21	50
Age 12	24	129	4	7	55	5	31	4	6	16	46
Age 13	15	111	2	5	54	5	24	4	5	8	38
Age 14	12	98	1	5	49	2	25	4	4	7	30
Age 15	10	87	1	5	48	1	23	4	1	6	23
Age 16	8	71	1	2	45	0	18	1	0	6	16
Age 17	7	63	0	1	43	0	15	1	0	6	12
Age 18	7	53	0	1	39	0	10	1	0	6	11
Age 19	6	44	0	1	32	0	8	1	0	5	7
Age 20	6	39	0	1	31	0	7	0	0	5	4
Totals	3,727	2,781	470	885	1,167	1,461	676	602	169	1,267	954

* Age measured relative to grant-year of the disclosed patent.
† This column reports the number of SSO patents disclosed at a given age.

Table 8: Age Effects for SSO and Control Patents

	Pooled Sample		ANSI		IEEE		IETF		ITU	
	SSO	Control	SSO	Control	SSO	Control	SSO	Control	SSO	Control
Age -2	0.709		0.050		0.476		1.038		0.082	
Age -1	1.412	0.683	1.019	0.593	1.294	0.639	1.428	0.601	1.154	0.617
Age 0	1.741	0.974	1.524	0.855	1.561	0.928	1.770	0.776	1.721	0.896
Age 1	1.939	1.139	1.869	1.020	1.825	1.093	1.808	0.934	1.838	1.038
Age 2	2.100	1.151	2.046	1.017	1.853	1.110	2.074	0.940	2.127	1.063
Age 3	2.110	1.106	2.023	0.950	1.952	1.061	2.204	0.925	2.029	1.020
Age 4	2.174	1.049	2.103	0.865	1.929	1.006	2.381	0.893	2.063	0.934
Age 5	2.096	0.953	2.074	0.758	1.970	0.935	2.223	0.800	1.925	0.835
Age 6	2.086	0.849	1.924	0.681	2.003	0.869	2.358	0.705	1.835	0.693
Age 7	2.017	0.767	1.889	0.629	1.918	0.828	2.223	0.659	1.781	0.565
Age 8	1.930	0.647	1.788	0.527	1.993	0.765	2.447	0.538	1.483	0.439
Age 9	1.875	0.569	1.747	0.467	1.967	0.642	2.409	0.415	1.407	0.340
Age 10	1.900	0.488	1.403	0.384	1.937	0.528	2.514	0.229	1.576	0.268
Age 11	1.890	0.394	1.776	0.244	1.936	0.394	2.218	0.272	1.545	0.247
Age 12	2.036	0.312	1.602	0.173	2.363	0.263	2.466	0.167	1.515	0.088

Regressions based on Equation (1), including a full set of unreported application- and citing-year effects.